

METHOD AND APPARATUS FOR ANALYZING A
MEDIA PATH IN A PACKET SWITCHED NETWORK

BACKGROUND

The Real Time Transport Protocol (RTP) is the primary protocol used for real time media transport over Internet Protocol (IP) networks. The RTP protocol is used for carrying voice, video, real time text, and other media types. In order to assess the reception quality of an RTP media stream, an associated Real Time Control Protocol (RTCP) feedback channel is used by the receiver to send information back to the sender.

There are a number of cases where it is undesirable to play out media. For example, a session setup protocol may wish to assess media channel viability before committing to forming a full media session and before alerting a user of a media call. It also may be undesirable to play out media when doing synthetic load testing or when there is a need to fill a channel to a constant packet rate to emulate a constant-bit-rate encoder or to foil traffic analysis attacks.

In one example, a first party may make an IP phone call to a second party. A signaling protocol is used to establish the phone call. The call signaling path may be operating correctly, thus allowing a signaling stream between the two parties. However, the media path is different from the signaling path. This means the correct operation of the signaling exchanges does not imply the ability to successfully exchange media packets.

Other communication networks, such as a Public Switched Telephone Network (PSTN) have the ability through mechanisms like Continuity Test (COT) to delay media play out while determining the viability of the media channel. A similar capability is desirable for Internet Protocol (IP) networks. It would be desirable to implement this capability with minimal or no new protocol machinery, and within the confines of existing IP standards.

Ideas have been proposed even in the packet domain for sending certain test packets for verifying that a media stream can be adequately supported over the IP network. Previous ideas include sending ping-pong packets or sending Named Signaling Events (NSEs) as described in RFC2833. Other ideas have suggested sending early media, or using "probe" packets before establishing the call. However, these ideas require awkward associations of RTP streams and non-standard behavior of endpoints.

The test packets may use an independent test path which presents several problems. First the test path and the media path may take different routes over the IP network. Another problem is that the test packets in the test path may be processed in the receiving device in a different way than the media stream received in the media path.

SUMMARY OF THE INVENTION

No-op media payload packets are used to analyze a media path in a packet switched network. In one embodiment, the no-op packets are Real Time Protocol (RTP) payload packets that contain no media content. A Real Time Control Protocol (RTCP) report is generated for the received RTP no-op packets. A marker bit is set in one or more of the no-op packets that triggers the no-op packet receiver to send back the RTCP report. The adequate operation of the media path is determined by examining the statistics in the RTCP report.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram showing how no-op payload packets are used for evaluating a media path.

FIG. 2 is a block diagram showing how the no-op packet evaluation scheme is negotiated between two endpoints.

FIG. 3 is a flow diagram showing how the no-op payload packets are generated.

FIG. 4 is a flow diagram showing how the media path is analyzed using the no-op media packets.

FIG. 5 is a flow diagram showing how the analysis information about the media path is used for establishing the media path.

FIG. 6 is a diagram of an RTP no-op packet.

FIG. 7 is a detailed diagram of a network device that processes the no-op packets.

DETAILED DESCRIPTION

A new media payload type is referred to as a no-op payload packet or simply a no-op packet. One application for the no-op packet is in the Real Time Transport Protocol (RTP),

5 which is specified in RFC3550. The RTP protocol is used for carrying a real time media stream over an IP network. The media stream can contain any type of real time media, including not just audio or video, but also other media such as real time text transmission, or voice-band data such as fax or modem information. The no-op payload packets, when used in RTP, are referred to as RTP no-op packets. The no-op packet in one embodiment has no
10 media payload and uses a one-bit opcode in the RTP payload for requesting a media path continuity report.

The no-op payload packets are literally a no-op coder. When negotiated through a Session Description Protocol (SDP – RFC2327), H.245, or some other media description protocol, any sender is able to send RTP packets with this coder. Any receiver who has
15 expressed willingness to receive this coder processes the RTP no-op packets according to the basic RTP specification and then discards the RTP no-op packets without performing actual media play-out.

The RTP no-op packets allow the instrumentation mechanisms in RTP, namely the information in Real Time Control Protocol (RTCP) reports, to be evaluated prior to actually
20 sending voice, audio, or other media in an RTP session. For example, an implementation may decide that the media stream is not sent for this session if the RTCP report indicates a low quality media path.

FIG. 1 shows a Wide Area Network (WAN) or Local Area Network (LAN) referred to generally as IP network 26 that is used for transferring a real time media stream 38. Plain
25 Old Telephone Service (POTS) phone 12, Private Branch Exchange (PBX) 11A, etc. are coupled to the IP network 26 through voice gateway 16. Another POTS phone 32, PBX 11B, etc. are coupled to the IP network 26 through voice gateway 30. Voice Over Internet Protocol (VoIP) phones 10A and 10B and computers 14 and 34 are coupled directly to the IP network 26.

30 Any of the voice gateways 16, 32, VoIP phones 10A, 10B, and computers 14 and 34 include any combination of hardware and software for receiving or sending a media stream 38. The media stream 38 can include any real time media such as audio data, video data, and real time text such as the text used for the hearing impaired.

Any combination of the phones 10A, 10B, 12, and 32; PBXs 11A and 11B; computers
35 14 and 34; and gateways 16 and 30 are referred to below generally as endpoints, users, senders or receivers, etc. One of the endpoints initiates a media call to another endpoint. If

5 one of the phones 12 or 32 initiates the call, the gateway 16 or 30, respectively, would make the VoIP call on behalf of the phone. For example, the phone 12 may make a VoIP phone call to phone 32. A caller dials a phone number using phone 12 and the gateway 16 then determines the IP address for the gateway 30 serving phone 32. The two gateways 16 and 30 then conduct a signaling protocol for establishing the VoIP call between phone 12 and phone 10 32.

In another scenario, the computer 14 may want to receive a video stream or audio stream from computer 34 which operates as a music and video content server. In another scenario, the computer 14 may wish to make a VoIP call to VoIP phone 10B. The computer 14 initiates the media connection to VoIP phone 10B. The computer 14 may establish the 15 media connection through a gateway or may access VoIP phone 10B through the IP network 26, without using a gateway. Media connections can also be initiated from any of the other devices 10A, 10B, 11A, 11B, 32, or 34.

In the example described below, the phone 12 initiates a VoIP phone call to phone 32. For clarity, operations will be described as being performed by the phones 12 and 32.

20 However, some of these operations could be performed by the gateways 16 and 30. The phone number for phone 32 is entered at phone 12. There may be an intermediary device, such as a call controller that determines the IP address for gateway 30 or phone 32 based on the phone number dialed by phone 12. During this initial signaling, the two gateways 16 and 30 agree to use the no-op payload packet verification process described below.

25 Prior to transmitting the media stream 38, gateway 16 sends one or more no-op payload packets 22. In this example, the no-op payload packets are RTP no-op packets and three of the RTP no-op packets are sent from gateway 16 to gateway 30. However any type of no-op payload packets 22 can be sent depending on the media stream carrier protocol. Any number of the no-op payload packets can be sent by gateway 16. However, the number 30 of no-op packets should be limited to prevent substantial delay in completion of the call setup.

The RTP no-op packets 22 have the same format as conventional RTP packets except that the packet payload contains no media. In the example shown, a marker bit 24 in the packet payload is not set in the first two RTP no-op packets 22 but the marker bit 24 is set in 35 the third RTP no-op packet 22. It should be understood that this is only one example and any number of no-op packets 22 can be sent. For example, in other implementations less than

5 three, or more than three, no-op packets are sent with the last no-op packet having a set marker bit 24. In another example, only one no-op packet is sent with that single no-op packet having a set marker bit.

When the gateway 30 receives the RTP No-Op packet with the marker bit set, the gateway 30 immediately generates an RTCP report 28 that includes Quality of Service (QoS) data for all of the RTP no-op packets 22 previously received. The RTCP report 28 is generated in the same manner used for generating RTCP reports for conventional RTP packets containing media payloads. For example, the phone 32 may count the number of received RTP no-op packets 22 as well as tracking latency and jitter statistics.

The RTCP report 28 would normally be sent by gateway 30 to gateway 16 about every 5 seconds. However, the set marker bit 24 in the RTP no-op packet 22 causes the gateway 30 to immediately send the RTCP report 28 to gateway 16. In this example, the RTCP report 28 would indicate that three RTP packets were received by gateway 30 if no packets were lost or unduly delayed. If the information in RTCP report 28 indicates the media path 18 carrying the RTP no-op packets 22 is operating adequately, signaling 36 is sent causing phone 32 to ring. This can be the un-modified "alert" signaling, or could be QoS-aware signaling using preconditions as described in RFC3312. Then the two phones 12 and 32 start sending the real time media stream 38.

If the RTCP report 28 indicates that the media path 18 is not operating adequately the signaling protocol is so informed, and can take action based on local or global policy. Possible policies include, but are not limited to: aborting the call such that the phone 32 does not ring, continuing with the call but making note to the accounting and billing system that the call may have to be discounted due to poor quality, or eliminating the problematical media stream if the call carries multiple media streams (e.g. audio and video).

In the case where the call is aborted, the aborted call may be established over an alternate path in IP network 26 or over an alternative medium, such as over a circuit switched network. A poorly performing media path may be identified when the RTCP report 28 indicates that fewer RTP no-op packets were received by gateway 30 than were sent by gateway 16. A media path problem may also be identified in the RTCP report 28 if the amount of latency or jitter in the RTP no-op packets 22 is outside of some expected range. In another scenario, a bad media path is indicated when the RTCP report 28 is not received by

5 gateway 16 within some amount of time after gateway 16 sends out an RTP no-op packet 22 with the set marker bit 24.

The no-op packet call verification process described above can be performed independently for each endpoint. For example, in a bidirectional transmission, such as with a phone call between phones 12 and 32, each phone can make a determination whether or not to establish the call based on the information in the RTCP report 28 received from the other
10 phone or gateway, or the lack of a RTCP report 28.

If the media transmission is unidirectional, then only one endpoint may need to verify the media channel. Unidirectional media transmissions are used for music or video streams transmitted unidirectionally from a media server to a user. In this application, only the path
15 from the media server to the user is important for media quality considerations. Thus, only the user receiving the media stream needs to send RTP no-op packets and evaluate the RTCP report 28 sent by the media server.

FIG. 2 shows generally how the initial call signaling is conducted between two endpoints. The Request For Comment (RFC) 3312 specifies one example of a real time
20 media signaling protocol. Of course, other types of real time media stream signaling protocols can also be used. Examples of signaling protocols include Session Initiation Protocol (SIP), Session Definition Protocol (SDP) offer-answer, H323, and Media Gateway Control Protocol (MGCP).

Signaling messages 42 are sent from phone 12 to phone 32 prior to transmitting the
25 media stream 38 shown in FIG. 1. The signaling messages 42 include a request or notification that the phone 12 intends to conduct a media path continuity check using RTP no-op packets. The signaling 42 may need to be conducted through a call controller 40. The phone 32 agrees to conduct a continuity check using the RTP no-op packets by sending acknowledgement signaling 44 back to phone 12.

30 The phone 32 can also send additional notification back to phone 12 in signaling 44 that indicates phone 32 also intends to send RTP no-op packet to phone 12 for continuity checking. The two phones 12 and 32 thus both agree to send and receive the RTP no-op packets 22 (FIG. 1) prior to sending and receiving a media stream.

35 RTP No-Op Packet Sender

5 FIG. 3 shows in further detail the sequence followed by the RTP no-op packet sender. In block 50, sequence numbers are assigned to the RTP no-op packets starting at the next sequence number for the RTP session. The sequence number for each no-op packet is incremented by one in accordance with the RTP specification. In block 52, a timestamp is inserted into the no-op packet according to a packet time (ptime) parameter, also in
10 accordance with RTP operation. The ptime parameter is conveyed via the signaling protocol, for example via SDP. It denotes a packetization interval identifying how many milliseconds of media are encoded into a single RTP packet.

 The sender identifies in block 54 the last RTP no-op packet to be sent. If it is not the last no-op payload packet to be sent, then the sender transmits the no-op packet in block 58.
15 If it is the last no-op packet, then the sender sets the marker bit in the RTP header of the no-op payload packet in block 56 causing the no-op packet receiver to immediately send back an RTCP report. The RTP no-op packet with the set marker bit is then transmitted to the receiver in block 58.

 The RTCP report is sent within the session RTCP bandwidth limits and other RTCP
20 sending rules. The scheme for sending the RTP no-op packets from the opposite end of the call are the same.

RTP No-Op Packet Receiver

 FIG. 4 shows in further detail the operations used at the processing device receiving
25 the RTP no-op packets. In block 60, a syntax check is performed on the received RTP no-op packets in the same way as any other RTP payload. In block 62, jitter buffer operations are performed on the RTP no-op packets in the same manner as any other RTP payload. The jitter buffer operations may include Synchronization Source field (SSRC), sequence, and timestamp processing.

30 In block 64, the RTCP statistics are updated for the RTP no-op packets in the same manner as for any other RTP payload format. If the marker bit in the RTP no-op packet is not set in block 66 or other RTCP sending rules are not met in block 67, the receiver continues to update RTCP statistics for any additionally received RTP no-op packets. If the marker bit is set for a received RTP no-op packet in block 66, other RTCP sending rules are
35 checked in block 67. For example, it is determined if the RTCP bandwidth constraints allow sending the RTCP report. If these RTCP sending rules are met in block 67, the RTCP report

5 is immediately sent to the no-op packet sender in block 68. The receiver rules for the other endpoint are the same,

Media Channel Acceptance

FIG. 5 describes the offer-answer exchange between the two endpoints. In block 70,
10 one of the endpoints proposes using the "no-op" coder and payload format during the session. In one example, an SDP (or H.245) with an offer-answer style exchange is used for conducting the RTP no-op notification.

In block 72, the notification may also optionally carry an RFC3312 precondition to ensure that no user is alerted of a call before the acceptance procedure using the no-op
15 payload packets is successfully completed. The preconditions could be used in either direction, although it may be typical to enforce their use in both directions. For example, when there is an RTP session in both directions of a full-duplex VOIP call, both directions would be required to pass the no-op packet acceptance test.

After completing the first RFC3312 offer-answer exchange, the sender starts sending
20 a configurable number of RTP no-op payload packets to the receiver in block 74. The marker bit is set in the last no-op payload packet in block 76. The receiver conducts the RTCP processing and returns the RTCP report after seeing the set marker bit. If the packet with the marker bit was lost, the RTCP report may be sent after the normal startup RTCP delay interval.

25 The sender waits for the RTCP report in block 78. If the RTCP report is not received in block 86 within some time period after sending out the last RTP no-op packet (i.e., timeout), the media path is declared unacceptable and the signaling protocol can decide what to do in block 88 according to any range of policies such as the ones described above. Alternatively, the procedure described above is repeated to cover the case of a single RTCP
30 report packet getting lost.

If the RTCP report is received in block 78 within an acceptable time period, the RTCP report is examined in block 82. Depending on the loss and jitter statistics, the receiver either declares the channel acceptable or unacceptable in block 82. If the media channel is deemed unacceptable, the signaling protocol can decide what to do in block 88 according to the
35 predetermined range of policies. If the channel is deemed acceptable, the precondition is marked as satisfied and the sender initiates (or responds to) the second phase of call

5 establishment in block 84 that allows the call to proceed. In particular, one of the endpoints may notify a user of the call by activating an annunciation device.

Referring to FIG. 6, the RTP no-op packet 22 is formatted to look like a normal RTP payload packet. This allows the no-op packet to operate and be processed in the same way as a voice, video or other type of RTP media packets. The RTP no-op packet will also take the
10 same path as the RTP packets that are sent after the call is set up. The only difference is that the RTP no-op packets 22, 23, 24 are not played out by the receiver.

The packet header 89 includes conventional RTP header content including a sequence number, timestamp value, synchronization source identifiers and contributing source (CSRC) identifiers. The timestamp value can be the real time-stamp used with actual media payloads
15 or the timestamp value can be simulated.

The RTP no-op payload 87 employs the marker bit 24 in the RTP header. If bit 0 is set, the receiver immediately sends out an RTCP report. Bits 1-31 in the payload 87 are reserved and in one example must all be set to zero. Additional padding bytes may be appended up to the negotiated packet time (ptime) value in SDP. Padding may be useful to
20 generate RTP packets that are the same size as other payload, such as a normal voice payload. Of course other formats can also be used in the RTP packet 22 for generating the no-op payload format.

FIG. 7 shows in further detail the elements in a network device 90 that sends or receives the RTP no-op packets. The device 90 can provide the functionality in any
25 combination of phones 12 and 32, computers 14 and 34, and gateways 16 or 30 shown in FIG. 1. For example, in one configuration the device 90 is either computer 14 or 34. The computer 14 or 34 performs all of the operations used for conducting the media session, including dialing another telephone or accessing another computer. In another configuration, the device 90 may be the gateway 16 or 30 shown in FIG. 1 and the media session is initiated
30 by a VoIP phone or computer that is connected to the gateway.

A processor 92 is coupled to a user interface 94 and a network interface 96. The network interface 96 connects to the IP network 26. The user interface 94 connects to any device needed to initiate the media connection. For example, user interface 94 could connect to a keyboard or to a phone.

35 A memory 98 contains executable software that contains both call signaling software 100 and RTP and RTCP software 102. The memory 98 represents any combination of

5 internal memory within the processor 92 or memory external to the processor 92. For example, the memory 98 can be any combination of flash, Read Only Memory (ROM), Random Access Memory (RAM), or disk memory used for containing and executing the signaling software 100 and the RTP and RTCP software 102.

10 The signaling software 100 is used for establishing the initial media path between the two endpoints. The signaling software 100 includes the RTP no-op negotiation described above. The RTP and RTCP software 102 conducts conventional RTP and RTCP with the addition of conducting media path verification using the RTP no-op packets and responding with RTCP reports as described above.

15 Thus, the RTP no-op payload format, with the special use of the RTP marker bit and RFC3312-style offer-answer and call signaling preconditions are used to perform channel acceptance testing as part of VoIP call establishment. Standard RTP packets are used rather than something special like a ping or an NSE. The marker bit is used along with RTCP to control the feedback loop.

20 This no-op payload scheme solves the problem of assessing media connectivity and channel quality in packet switched networks. The no-op scheme can be used in any and all VOIP or media endpoints, but particularly those for sale to and through service providers accustomed to channel assessment capabilities, such as those provided in circuit switched networks. The no-op payload format can also be used for other applications, such as synthetic load testing and to foil traffic analysis.

25 The system described above can use dedicated processor systems, micro controllers, programmable logic devices, or microprocessors that perform some or all of the operations. Some of the operations described above may be implemented in software and other operations may be implemented in hardware.

30 For the sake of convenience, the operations are described as various interconnected functional blocks or distinct software modules. This is not necessary, however, and there may be cases where these functional blocks or modules are equivalently aggregated into a single logic device, program or operation with unclear boundaries. In any event, the functional blocks and software modules or features of the flexible interface can be implemented by themselves, or in combination with other operations in either hardware or
35 software.

5 Having described and illustrated the principles of the invention in a preferred embodiment thereof, it should be apparent that the invention may be modified in arrangement and detail without departing from such principles. We claim all modifications and variation coming within the spirit and scope of the following claims.